

CLAIMS

1 1. A method of accelerating the operation of a load balancer by an accelerator switch,
2 comprising:

3 receiving, by the accelerator switch, packets directed to the load balancer;

4 determining, for at least one of the received packets, whether the packets match an
5 entry of a list of packet groups, by comparing fewer than five parameters of the packets to
6 respective fields of entries of the list; and

7 forwarding, by the accelerator switch, at least one of the received packets, directly to its
8 destination, responsive to the determining.

1 2. A method according to claim 1, wherein determining whether the packets match an
2 entry of the list comprises comparing three or fewer parameters of the packets to respective
3 fields in the list.

1 3. A method according to claim 2, wherein determining whether the packets match an
2 entry of the list comprises comparing two parameters of the packet to a respective field in the
3 list.

1 4. A method according to claim 2, wherein determining whether the packets match an
2 entry of the list comprises comparing a single parameter of the packet to a respective field in
3 the list.

1 5. A method according to claim 1, wherein receiving packets directed to the load balancer
2 comprises receiving packets directed from a client to a Web site associated with the load
3 balancer and forwarding at least one of the received packets directly to its destination
4 comprises forwarding the packets from the clients to one of the servers of the Web site without
5 passing through the load balancer.

1 6. A method according to claim 5, wherein determining whether the packets match an
2 entry of the list comprises comparing the source IP address and source port of the packets to
3 respective fields in the list.

1 7. A method according to claim 5, wherein the compared parameters do not include a
2 destination address.

1 8. A method according to claim 1, wherein receiving packets directed to the load balancer
2 comprises receiving packets directed from a server to a client and forwarding at least one of the
3 received packets directly to its destination comprises forwarding the packets from the server to
4 the client without passing through the load balancer.

1 9. A method according to claim 8, wherein determining whether the packets match an
2 entry of the list comprises comparing the destination IP address and destination port of the
3 packets to respective fields in the list.

1 10. A method according to claim 8, wherein the compared parameters do not include a
2 source address.

1 11. A method according to claim 1, wherein forwarding at least one of the received packets
2 comprises forwarding packets for which a matching entry was found.

1 12. A method according to claim 1, wherein the load balancer operates in a half NAT or full
2 NAT mode.

1 13. A method of creating an entry in a list which correlates between packet groups and
2 respective destination servers, comprising:

3 receiving, by an accelerator, a packet directed from or to a load balancer which changes
4 at least the destination IP address of packets it forwards to servers; and

5 creating, by the accelerator, an entry in the list of destination servers, responsive to the
6 received packet.

1 14. A method according to claim 13, wherein creating the entry comprises creating an entry
2 which does not include a destination address of a Web site.

1 15. A method according to claim 13, wherein the packet is directed from or to a load
2 balancer operating in a half NAT mode.

1 16. A method according to claim 13, wherein the packet is directed from or to a load
2 balancer operating in a full NAT mode.

1 17. A method according to claim 13, wherein receiving the packet comprises receiving a
2 packet directed from the load balancer to a server.

1 18. A method according to claim 13, wherein receiving the packet comprises receiving a
2 packet directed from a server to the load balancer.

1 19. A method according to claim 13, wherein creating the entry comprises creating the
2 entry using substantially only information in the received packet as it was received.

1 20. A method according to claim 13, wherein creating the entry comprises creating the
2 entry using information not included in the received packet as it was received.

1 21. A method according to claim 20, wherein creating the entry comprises creating the
2 entry using information from a copy of the received packet, previously received by the
3 accelerator.

1 22. A method according to claim 21, wherein receiving the packet comprises receiving a
2 packet from the load balancer and creating the entry comprises creating the entry using
3 information from the received packet and from a copy of the received packet forwarded to the
4 load balancer.

1 23. A method according to claim 13, comprising receiving, by the accelerator, a packet
2 directed from or to an additional load balancer and creating, by the accelerator, an entry in the
3 list of destination servers, responsive to the received packet.

1 24. A method according to claim 13, comprising:

2 receiving, by the accelerator, packets directed to a Web site handled by the load
3 balancer;

4 storing identification information and values of one or more parameters of the packets
5 directed to the Web site, in a temporary storage; and

6 searching the temporary storage for an entry which matches the received packet from
7 the load balancer,

8 wherein creating the entry in the list of destination servers of packet groups is
9 performed only if a match is found.

1 25. A method according to claim 24, wherein storing the identification information
2 comprises storing a unique identification number tagged to the packet by the accelerator.

1 26. A method according to claim 24, wherein storing the identification information
2 comprises storing at least one of the sequence and acknowledge fields of TCP packets.

1 27. A method according to claim 26, wherein storing the identification information
2 comprises storing a leading segment of the payload of the packet.

1 28. A load balancing accelerator, comprising:

2 an input interface which receives packets directed to a load balancer;

3 a table which lists packet groups and their respective destination servers, the table
4 having physical entries which can accommodate different field sets for storage of data entries;

5 a comparator which compares at least one of the packets directed to the load balancer
6 to one or more of the data entries of the table;

7 a forwarding unit which forwards at least one of the packets for which a match was
8 found by the comparator, directly to a server, responsive to the contents of the matching data
9 entry; and

10 a controller which determines in which field set, from the plurality of different field sets,
11 each of the data entries of the table is stored.

12 29. An accelerator according to claim 28, wherein the controller comprises a user interface
13 through which a user may configure the field sets in which the data entries of the table are
14 stored.

15 30. An accelerator according to claim 28, wherein the controller automatically determines
16 the field sets in which the data entries are stored.

17 31. An accelerator according to claim 30, wherein the controller transmits one or more
18 packets to the load balancer and examines the response of the load balancer to determine the
19 field sets in which the data entries are stored.

20 32. An accelerator according to claim 28, wherein the controller determines the field sets in
21 which the data entries of the table are stored, such that at a single time all the data entries of the
22 table are stored in the same field sets.

23 33. An accelerator according to claim 28, wherein the controller determines the field sets in
24 which the data entries of the table are stored, such that at least during some periods of
25 operation of the accelerator, the table includes at least two data entries stored in different field
26 sets.

27 34. An accelerator according to claim 28, wherein at least one of the physical entries of the
28 table can be configured for use with different field sets.

1 35. An accelerator according to claim 28, wherein the table comprises a plurality of sub-
2 tables with physical entries having different field sets.

1 36. An accelerator according to claim 28, wherein the input interface receives packets
2 directed to a plurality of load balancers and wherein the data entries corresponding to a first
3 load balancer are stored in a first set of fields and data entries corresponding to a second load
4 balancer are stored in a second set of fields different from the first set of fields.

1 37. A load balancing accelerator, comprising:

2 an input interface which receives packets directed to a load balancer;

3 a table which lists packet groups and their respective destination servers;

4 a comparator which compares at least one of the packets directed to the load balancer
5 to at least one of the entries of the table;

6 a forwarding unit which forwards directly to a server, at least one of the packets for
7 which a match was found by the comparator, responsive to the contents of the matching entry,
8 the forwarding unit being capable of operating in a plurality of operation modes, at least one of
9 the operation modes including changing at least one of the fields of the forwarded packets; and
10 a controller which determines in which mode the forwarding unit operates.

1 38. An accelerator according to claim 37, wherein the forwarding unit is capable of
2 performing splicing.

1 39. An accelerator according to claim 37, wherein the controller determines the operation
2 mode of the forwarding unit based on a user configuration.

1 40. An accelerator according to claim 37, wherein the controller determines the operation
2 mode of the forwarding unit based on the contents of packets directed from or to the load
3 balancer.

1 41. An accelerator according to claim 37, wherein the controller determines the operation
2 mode of the forwarding unit by comparing the contents of packets from the load balancer with
3 packets directed to the load balancer.